A Generalization of Pratt-Arrow Measure to Nonexpected-Utility Preferences and Inseparable Probability and Utility

Robert F. Nau

Fuqua School of Business, Duke University, Durham, North Carolina 27708-0120
robert.nau@duke.edu

The Pratt-Arrow measure of local risk aversion is generalized for the n-dimensional state-preference model of choice under uncertainty in which the decision maker may have inseparable subjective probabilities and utilities, unobservable stochastic prior wealth, and/or smooth nonexpected-utility preferences. Local risk aversion is measured by the matrix of derivatives of the decision maker's risk-neutral probabilities, without reference to true subjective probabilities or riskless wealth positions, and comparative risk aversion is measured without requiring agreement on true probabilities. Risk-neutral probabilities and their derivatives are shown to be sufficient statistics for approximately optimal investment and financing decisions in complete markets for contingent claims.
(*Criteria for Decision Making Under Risk and Uncertainty; Risk Aversion; Uncertainty Aversion; Expected-Utility Theory; Nonexpected-Utility; Smooth Preferences*)

1. Introduction

Risk aversion is commonly defined as a departure from expected-value-maximizing behavior: A risk-averse person always prefers a riskless wealth position to a risky position with the same expected value. For a decision maker with expected-utility preferences, determinate probabilities, state-independent utility, and observable wealth, this definition is as good as any other, and the decision maker's degree of risk aversion is conveniently represented by the Pratt-Arrow measure (de Finetti 1952, Pratt 1964, Arrow 1965), which quantifies the local curvature of her Bernoulli utility function. However, in many economically interesting situations involving uncertainty, the definition of risk aversion by reference to expected values and riskless wealth can be problematic, especially from the viewpoint of an observer. Probabilities may be subjective and utilities may be state dependent, in which case probabilities and expected values are not uniquely revealed by preferences. If the individual has nonexpected-utility preferences—e.g., if she exhibits aversion to uncertainty—her beliefs may not even be representable by additive probabilities. Individuals may also have significant unobserved prior stakes in events and they may face uninsurable risks, in which case riskless wealth positions may be ill defined or unattainable, and without independent knowledge of correlations with prior wealth it is not clear whether the acquisition of another "risky" asset produces an increase or decrease in overall risk. Yaari (1969) has suggested a more elementary definition of risk aversion that does not depend on observability of probabilities or prior wealth, namely that a decision maker is risk averse if her preferences are *payoff convex*, that is, convex with respect to deterministic mixtures of payoffs within states of the world. The latter definition, which we shall adopt here, agrees with the conventional one for decision makers who are expected-utility maximizers, and it also applies in a straightforward way to decision makers with

Figure 1   Indifference Curves of Four Individuals for Wealth in Two States
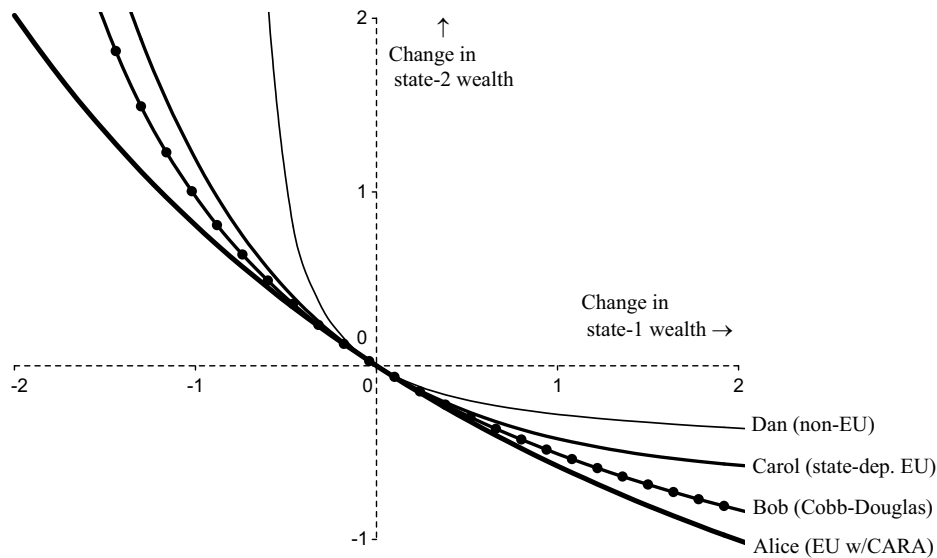
state-dependent utility or nonexpected-utility prefer-ences,[1] although it is not the only possible definition of risk aversion in the latter cases (e.g., Machina 1995, Epstein 1999). This paper derives a matrix-valued measure of risk and uncertainty aversion that gener-alizes the Pratt-Arrow measure to the broader frame-work of Yaari's definition.

To motivate the discussion, Figure 1 shows the indifference curves of four hypothetical individuals with respect to distributions of wealth over two states of the world. The curves are aligned so that the ori-gin of coordinates corresponds to the status quo for each individual, although it is not a riskless posi-tion for anyone: They all have stochastic prior wealth. Alice claims to be an expected-utility maximizer having the utility function $U(\mathbf{w}) = -p\exp(-rw_1) - (1-p)\exp(-rw_2)$, with $p = 1/3$, $r = 0.2$, and prior wealth distribution $(w_1, w_2) = (2.8, 4.3)$. Thus, Alice exhibits constant absolute aversion to risk with a risk aversion coefficient of 0.2, and she evidently assigns probability 1/3 to State 1. Bob has a Cobb-Douglas utility function, $U(\mathbf{w}) = w_1^\alpha w_2^{1-\alpha}$, with $\alpha = 0.355$ and

prior wealth distribution $(2.5, 3.0)$. Of course, this is equivalent to having expected-utility preferences with logarithmic utility and a probability of 0.355 for State 1, but Bob insists that he does not subscribe to expected-utility theory and has no opinion con-cerning the probabilities of the states; his preferences just "are what they are." Carol has state-dependent expected-utility preferences represented by the utility function $U(\mathbf{w}) = -p\exp(-r_1 w_1) - (1-p)\beta\exp(-r_2 w_2)$ with $r_1 = 1$, $r_2 = 0.05$. Thus, Carol is much more risk averse than Alice in State 1 (as measured by concav-ity of utility for State-1 wealth), while she is less risk averse in State 2. However, she is evasive about her other utility parameters: She says it's possible that her probability of State 1 is $p = 1/3$, her rate of util-ity substitution between states is $\beta = 1$, and her prior wealth is $(2.8, 2.0)$. However, it could be that $p = 1/2$ and $\beta = 2$. Or, perhaps her prior wealth is $(3.023, 2.0)$ while $p = 1/3$ and $\beta = 0.8$. As she is fond of point-ing out, it really doesn't matter: Her preferences for changes in wealth would be the same in all three cases. Carol also suspects that Alice has been less than truthful: She thinks that Alice has additional, undis-closed wealth in State 1 and that Alice's true probabil-ity for State 1 is really greater than 1/3, but she can't prove it merely by watching how Alice bets or trades with others.

---

[1] Risk aversion concepts for nonexpected-utility preferences under *risk* are discussed by Chew et al. (1987), Chew (1995), Cohen (1995), and Courtault and Gayant (1998). Here the focus is on choice under *uncertainty*, where probabilities of events are at best subjective and at worst indeterminate.

Finally, Dan has nonexpected-utility preferences represented by the utility function $U(\mathbf{w}) = -\exp(-r(pw_1 + (1-p)w_2)) - \exp(-r(qw_1 + (1-q)w_2))$, with $p = 0.05$, $q = 0.95$, $r = 2$, and a prior wealth distribution of (1.25, 1.0).[2] Coincidentally, everyone has the same marginal rate of substitution for wealth between States 1 and 2, namely that $3 in State 1 is equivalent to $2 in State 2—in other words, the *risk-neutral probability* of State 1 is 0.4 for everyone— perhaps because they have already engaged in trade with each other, or perhaps because they are in contact with an external market in which the price of a State-1 Arrow security is 0.4.

Suppose that we wish to measure and compare the degrees of local risk aversion of these four individuals to determine their relative propensity to gamble, purchase risky assets, or take other decisions under uncertainty, starting from their status quo wealth positions rather than primordial riskless positions. The usual definitions of risk aversion and risk premia do not apply because probabilities and current wealth positions are ill defined or unobservable, and only Alice's local risk preferences can be characterized by a scalar Pratt-Arrow measure. However, all four are plainly risk averse in the sense of having preferences that are payoff convex, and they can even be strictly ordered in terms of their degrees of local risk aversion: Alice < Bob < Carol < Dan. The indifference curves passing through their status quo wealth positions are the frontiers of their respective sets of acceptable gambles—which are observable—and following Yaari (1969), an individual with a strictly smaller set of acceptable gambles is "more risk averse." To compare local risk preferences in these terms, it is unnecessary to know anyone's probabilities, utility functions for money, or prior wealth: It suffices to consider the slope and curvature of their respective indifference curves. Yaari explored the two-dimensional

case, showing that the degree of local risk aversion could be measured by the second derivative of the parameterized indifference curve. However, the latter measure does not generalize easily to higher dimensions.

This paper derives a measure of local risk aversion for the general n-dimensional case, in which the curvature of indifference curves is measured by the matrix of derivatives of the decision maker's risk-neutral probabilities—i.e., the derivatives of the normalized gradient of her ordinal utility function. In the special case where preferences are separable across mutually exclusive events, the new risk aversion measure is additively separable and reduces to a state-dependent form of the Pratt-Arrow measure, in which case the decision maker is uncertainty neutral.[3] For such an individual, the risk-neutral distribution is the "correct" probability distribution to use in conjunction with the Pratt-Arrow measure when calculating risk premia in the presence of background risk and/or state-dependent utility. Under more general preferences, the new measure incorporates both second-order risk aversion and second-order uncertainty aversion. A nonneutral attitude toward uncertainty is revealed when a decision maker is uniformly more risk averse toward bets on some events than toward others, as discussed in a companion paper (Nau 2002).

The organization of this paper is as follows. Section 2 presents the modeling framework and definitions of risk aversion and risk premia. Section 3 derives the main result, namely a generalization of the Pratt-Arrow measure for the general n-state model. Section 4 considers the special case of separable preferences—essentially, state-dependent subjective expected utility without uniquely determined probabilities—for which the risk aversion measure is

---

[2] A utility function of this form can be used to model uncertainty aversion and rationalize the Ellsberg and Allais paradoxes (Nau 2002, Klibanoff et al. 2002). Here, it is as if Dan has linear utility for money but he is uncertain about the probability of State 1, which he feels is equally likely to be either 0.05 or 0.95, and he has "constant absolute aversion to uncertainty" with an uncertainty aversion coefficient of 2.

[3] A decision maker is "uncertainty averse" if she dislikes betting on events with ambiguous probabilities, as in Ellsberg's paradox, and "uncertainty neutral" otherwise. Second-order aversion to risk or uncertainty arises from the curvature of indifference curves, whereas first-order aversion arises from kinks in indifference curves (Segal and Spivak 1990). In the Choquet and maxmin expected-utility models, uncertainty aversion is exclusively a first-order effect, whereas the focus of this paper is on second-order effects—i.e., smooth rather than kinked preferences.

vector valued. Section 5 discusses comparative risk aversion and marginal investment behavior; §6 considers decision making in markets under uncertainty; §7 applies the market results to a project valuation decision, and §8 presents concluding comments.

2. The Model

The analytic framework used here is that of state-preference theory (Debreu 1959, Arrow 1964, Hirshleifer 1965), which includes, as special cases, models of expected utility, subjective expected utility, state-dependent utility, and Choquet and maxmin expected utility as they apply to monetary acts. In the state-preference framework, objects of choice are distributions of monetary wealth over states of the world, represented by vectors in $\Re^n$. The decision maker's preferences among such wealth distributions are minimally assumed to satisfy the basic axioms of consumer theory (completeness, transitivity, continuity), which imply that they can be represented by an ordinal utility function U and visualized in terms of indifference curves in payoff space, as illustrated in Figure 1. In general, U is determined only up to monotonic transformations and is therefore unobservable.

A decision maker who is risk neutral has linear indifference curves in payoff space, while one who is risk averse does not. The conventional definition of risk aversion, namely that a risk-averse individual always prefers a riskless wealth position to a risky position with the same expected value, has the following geometric interpretation: An indifference curve drawn through a point on the "45-degree certainty line" in payoff space must lie on or above the iso-expected-value line through the same point. A stronger definition, drawing on Rothschild and Stiglitz's (1970) concept of increasing risk, is that a risk-averse individual dislikes mean-preserving spreads in payoff distributions, which means that any movement along an iso-expected-value line away from the 45-degree certainty line (i.e., in a direction that takes payoffs in all states farther from the expected value) is dispreferred to the status quo, even if the status quo is already risky. A third definition of risk aversion, due to Yaari (1969), is that a

risk-averse individual has preferences that are payoff convex,[4] which implies that her ordinal utility function U is quasi-concave. The three definitions are equivalent for decision makers with expected-utility preferences,[5] but they differ for decision makers with nonexpected-utility preferences. Although the mean-preserving-spread definition characterizes a form of "local" risk aversion, in the sense that the definition has local implications even at risky wealth positions, it nevertheless admits local behavior that is arguably risk loving. In particular, it admits the possibility that for some risky wealth position $\mathbf{w}$ and finite gamble $\mathbf{z}$, both $\mathbf{w}+\mathbf{z}$ and $\mathbf{w}-\mathbf{z}$ are strictly preferred to $\mathbf{w}$. (For an illustration, see Machina 1995; see also Karni 1995.) Thus, for example, the decision maker might be willing to pay a fee for the privilege of placing a bet on an event and also (alternatively) willing to pay a fee for the privilege of placing the very opposite bet. This sort of behavior, which resembles gambling for its own sake, is expressly forbidden by Yaari's definition. On the other hand, the mean-preserving-spread definition requires prior wealth to be observable so that the certainty line can be located, and it requires the decision maker to be "probabilistically sophisticated" (Machina and Schmeidler 1992) so that expected values are uniquely determined, whereas Yaari's method makes no such requirements. For these reasons, Yaari's definition of risk aversion as payoff convexity of preferences (or equivalently, as quasiconcavity of utility) will be adopted henceforth.

Assume that preferences over wealth distributions are smooth, so that the utility function U that represents them is twice differentiable.[6] By virtue of monotonicity, the gradient of U at $\mathbf{w}$ is a nonnegative

---

[4] The preference relation "$\geq$" is [strictly] payoff-convex if $\mathbf{x} \geq \mathbf{z}$ and $\mathbf{y} \geq \mathbf{z}$ imply $\alpha\mathbf{x}+(1-\alpha)\mathbf{y} \geq [>]\mathbf{z}$ for $\alpha \in (0,1)$.

[5] For a decision maker with expected-utility preferences, local risk aversion at riskless wealth positions is sufficient to ensure concavity of the Bernoulli utility function (the univariate utility function for money), which implies that preferences are everywhere locally risk averse and also payoff convex.

[6] The smoothness assumption rules out Choquet or maxmin expected-utility preferences, but it permits nonexpected utility preferences that are arbitrarily close to Choquet expected utility in the fashion of Dan's preferences in Figure 1.

vector and it can be normalized to yield a probability distribution:

$$\boldsymbol{\pi}(\mathbf{w}) \equiv \frac{\nabla U(\mathbf{w})}{\left\|\nabla U(\mathbf{w})\right\|_1} = \frac{\nabla U(\mathbf{w})}{\mathbf{1} \cdot \nabla U(\mathbf{w})}.$$

$\boldsymbol{\pi}(\mathbf{w})$ is invariant to monotonic transformations of U and is observable. It is commonly known as a *risk-neutral probability distribution* because the decision maker prices very small assets in a seemingly risk-neutral manner with respect to it. More precisely, let $\mathbf{z}$ denote the payoff vector of a risky asset and let $P(\mathbf{z}; \mathbf{w})$ denote the *marginal price* that the decision maker is willing to pay for $\mathbf{z}$ at wealth $\mathbf{w}$, in the sense that she is willing to pay $\alpha P(\mathbf{z}; \mathbf{w})$ to receive $\alpha \mathbf{z}$ in the limit as α goes to zero. Then, $P(\mathbf{z}; \mathbf{w})$ is determined by

$$\lim_{\alpha \to 0}(U(\mathbf{w} + \alpha\mathbf{z} - \alpha P(\mathbf{z}; \mathbf{w})) - U(\mathbf{w}))/\alpha = 0,$$

for which the first-order condition is

$$P(\mathbf{z}; \mathbf{w}) = \mathbf{z} \cdot \boldsymbol{\pi}(\mathbf{w}) \equiv \mathrm{E}_{\boldsymbol{\pi}(\mathbf{w})}[\mathbf{z}].$$

Hence, the marginal price is the *risk-neutral expectation* of the asset under the local risk-neutral distribution $\boldsymbol{\pi}(\mathbf{w})$. $\mathbf{z}$ will be said to be a *neutral asset* at the current wealth position if $\mathrm{E}_{\boldsymbol{\pi}(\mathbf{w})}[\mathbf{z}] = 0$. The marginal price of $\mathbf{z}$ is the per-unit price at which the decision maker would buy or sell an infinitesimal share. Her *buying price* for $\mathbf{z}$ in its entirety, denoted $B(\mathbf{z}; \mathbf{w})$, is determined by

$$U(\mathbf{w} + \mathbf{z} - B(\mathbf{z}; \mathbf{w})) - U(\mathbf{w}) = 0,$$

while her *selling price* $C(\mathbf{z}; \mathbf{w})$, otherwise known as her *certainty equivalent* for $\mathbf{z}$, satisfies

$$U(\mathbf{w} + \mathbf{z}) - U(\mathbf{w} + C(\mathbf{z}; \mathbf{w})) = 0.$$

The buying and selling prices are generally similar, but not identical, as illustrated in Figure 2, and they are related by $C(\mathbf{z}; \mathbf{w}) = -B(-\mathbf{z}; \mathbf{w} + \mathbf{z})$.

The functional dependence of $\boldsymbol{\pi}(\mathbf{w})$ on $\mathbf{w}$ reveals the decision maker's attitude toward risk and uncertainty. If the decision maker is risk neutral (i.e., if U is both quasi-concave and quasi-convex), then she has linear indifference curves, $\boldsymbol{\pi}(\mathbf{w})$ is constant, and $P(\mathbf{z}; \mathbf{w}) = B(\mathbf{z}; \mathbf{w}) = C(\mathbf{z}; \mathbf{w})$ at all $\mathbf{w}$. If she is risk averse, $\boldsymbol{\pi}(\mathbf{w})$ varies with $\mathbf{w}$ according to the local

curvature of the indifference curves. Intuitively, a decision maker who is risk averse has diminishing marginal utility for all risky assets, hence her buying and selling prices for an asset will typically be less than its marginal price. To make this notion precise, let the *buying risk premium* associated with $\mathbf{z}$ at wealth $\mathbf{w}$, here denoted $b(\mathbf{z}; \mathbf{w})$, be defined as the difference between the asset's marginal price and its buying price

$$b(\mathbf{z}; \mathbf{w}) = \mathrm{E}_{\boldsymbol{\pi}(\mathbf{w})}[\mathbf{z}] - B(\mathbf{z}; \mathbf{w}).$$

The *selling risk premium*[7] is similarly defined by

$$c(\mathbf{z}; \mathbf{w}) = \mathrm{E}_{\boldsymbol{\pi}(\mathbf{w})}[\mathbf{z}] - C(\mathbf{z}; \mathbf{w}).$$

(Refer again to Figure 2.) Pratt's (1964) risk premium is the special case of the selling risk premium that obtains when U has the expected-utility representation $U(\mathbf{w}) = \sum_{j=1}^{n} p_j u(w_j)$, where u is a twice-differentiable state-independent utility function for money, $\mathbf{p}$ is a known probability distribution, and the decision maker begins in a state of riskless wealth $\mathbf{w} = x\mathbf{1}$, which is a point on the 45-degree certainty line in payoff space.

If the decision maker is an (possibly state-dependent) expected-utility maximizer, her risk-neutral probabilities are simply the product of her true probabilities and relative marginal utilities for money. That is, $\boldsymbol{\pi}_j(\mathbf{w}) \propto p_j u'_j(w_j)$, where u'_j is the first derivative of the utility function for money in state j (Drèze 1970). Under Pratt's assumptions of state-independent utility and riskless prior wealth, true probabilities and risk-neutral probabilities happen to coincide, and the selling risk premium can be interpreted as the amount of expected value the decision maker would give up to eliminate all risk following the involuntary acquisition of $\mathbf{z}$. The familiar result is that if $\mathbf{z}$ is small and "actuarially" neutral ($\mathrm{E}_{\mathbf{p}}[\mathbf{z}] = 0$),

[7] Buying and selling risk premia are called "compensating" and "equivalent" risk premia by Kimball (1990), who points out that they are essentially equivalent for small risks under state-independent expected utility and riskless prior wealth. In the context of state-dependent utility, Karni (1983, 1985) uses a generalization of the selling premium, while Kelsey and Nordquist (1991) prefer the buying premium, noting that the selling premium creates technical problems under some state-dependent utility functions.

Figure 2   Construction of Buying and Selling Prices and Risk Premia

Notes. The vector $\pi(\mathbf{w})$ is the risk-neutral distribution, i.e., the normalized gradient of the utility function U, at wealth $\mathbf{w}$. The dashed line is the tangent hyperplane (the subspace of neutral assets) at $\mathbf{w}$, whose normal vector is $\pi(\mathbf{w})$. The marginal price of an asset at $\mathbf{w}$ is its inner product with $\pi(\mathbf{w})$. Asset $\mathbf{z}$ is neutral at wealth $\mathbf{w}$, because its marginal price is zero, while $\mathbf{z}'$ is not. $\mathbf{w} + \mathbf{z}$ and $\mathbf{w} + \mathbf{z}'$ lie on the same indifference curve, below that of $\mathbf{w}$. The small dashed triangles are 45-degree right triangles with altitudes a, b, c, d. The prices and risk premia of $\mathbf{z}$ and $\mathbf{z}'$ at wealth $\mathbf{w}$ are:

Marginal Price	Buying Price	Selling Price	Buying Risk Premium	Selling Risk Premium
$P(\mathbf{z}; \mathbf{w}) = 0$	$B(\mathbf{z}; \mathbf{w}) = -a$	$C(\mathbf{z}; \mathbf{w}) = -b$	$b(\mathbf{z}; \mathbf{w}) = a$	$c(\mathbf{z}; \mathbf{w}) = b$
$P(\mathbf{z}'; \mathbf{w}) = -c$	$B(\mathbf{z}'; \mathbf{w}) = -(c+d)$	$C(\mathbf{z}'; \mathbf{w}) = -b$	$b(\mathbf{z}'; \mathbf{w}) = d$	$c(\mathbf{z}'; \mathbf{w}) = b - c$

then the selling risk premium is approximately

$$c(\mathbf{z}; x\mathbf{1}) \approx \tfrac{1}{2} r(x) \mathrm{E}_{\mathbf{p}}[\mathbf{z}^2] = \tfrac{1}{2} r(x) \mathrm{Var}_{\mathbf{p}}[\mathbf{z}]$$
$$= \tfrac{1}{2} \mathrm{Cov}_{\mathbf{p}}[\mathbf{z}, r(x)\mathbf{z}], \quad (1)$$

where $\mathbf{p} = \pi(x\mathbf{1})$ is the true distribution that coincides with the risk-neutral distribution and $r(x) = -u''(x)/u'(x)$ is the Pratt-Arrow measure of absolute risk aversion. (The buying risk premium converges to the same limit when $\mathbf{z}$ is small enough for (1) to be accurate, as will be seen.) Thus, $\tfrac{1}{2} r(x)$ is the decision maker's local price of risk, where risk is measured in terms of the variance of the asset, and the selling risk premium is nonnegative at constant wealth level x if and only if $r(x)$ is nonnegative. The function $r(x)$ also permits comparative risk aversion to be characterized in a simple way: If $r_1(x)$ and $r_2(x)$ are the risk aversion measures of Agents 1 and 2, respectively, then

Agent 1 is as risk averse as Agent 2 (in the sense of assigning greater or equal risk premia) if $r_1(x) \geq r_2(x)$ for all x, and Pratt and Arrow showed that in this case Agent 1 will also invest less than Agent 2 in a risky asset when given a choice between a single risky asset and a safe asset.

The buying risk premium rather than the selling risk premium will be used henceforth as the yardstick for measuring local risk aversion, for several reasons. First, the question that the selling premium is designed to address, namely how much (expected) value the decision maker would give up to eliminate all the risk she currently faces, is moot in the present setting of unobservable stochastic prior wealth and state-dependent or nonexpected-utility preferences. It is more natural to ask how much additional (marginal) value the decision maker would require as compensation for taking on a new risk.

Second, the buying risk premium has the convenient property that $b(\mathbf{z}+x;\mathbf{w}) = b(\mathbf{z};\mathbf{w})$ for any constant x, which is not true for the selling risk premium except in special cases. Third, and most importantly, the buying risk premium is a more natural indicator of risk aversion in the general state-preference framework because it directly measures the local quasi-concavity of utility. In particular,

PROPOSITION 1. *The decision maker is risk averse if and only if her buying risk premium is nonnegative for every asset at every wealth distribution.*

PROOF. Nonnegativity of the buying risk premium is essentially a definition of quasi-concavity: A function U is quasi-concave if and only if $\nabla U(\mathbf{w}) \cdot (\mathbf{w}' - \mathbf{w}) \geq 0$ whenever $U(\mathbf{w}') \geq U(\mathbf{w})$ (e.g., Theorem M.C.3 in Mas-Colell et al. 1995), and if U is monotonic, it suffices for this to hold when $U(\mathbf{w}') = U(\mathbf{w})$. Letting $\mathbf{w}' = \mathbf{w} + \mathbf{z} - B(\mathbf{z};\mathbf{w})$, we have $\nabla U(\mathbf{w}) \cdot (\mathbf{z} - B(\mathbf{z};\mathbf{w})) \geq 0$, which is equivalent to $\mathrm{E}_{\pi(\mathbf{w})}[\mathbf{z}] \geq B(\mathbf{z};\mathbf{w})$, which in turn is equivalent to $b(\mathbf{z};\mathbf{w}) \geq 0$. $\square$

3. A General Measure of Risk Aversion

The objective of this section is to characterize risk aversion in terms of second-order properties of preferences. As is well known, U is quasi-concave, and hence the decision maker is risk averse by our definition, if and only if at every $\mathbf{w}$ the Hessian matrix $D^2 U(\mathbf{w})$ is negative semidefinite in the subspace of neutral assets. This is not immediately helpful or economically significant, however, because $D^2 U(\mathbf{w})$ is not observable—that is, it is not uniquely determined by preferences. The observable second-order information resides instead in the matrix of derivatives of the risk-neutral probabilities, $D\pi(\mathbf{w})$, whose jkth element is

$$D\pi_{jk}(\mathbf{w}) = \partial \pi_j(\mathbf{w})/\partial w_k$$
$$= \frac{D^2 U_{jk}(\mathbf{w}) - \pi_j(\mathbf{w}) \sum_{h=1}^{n} D^2 U_{hk}(\mathbf{w})}{\mathbf{1} \cdot \nabla U(\mathbf{w})}, \quad (2)$$

where $D^2 U_{jk}(\mathbf{w})$ denotes $\partial^2 U(\mathbf{w})/\partial w_j w_k$. In principle, the elements of $D\pi(\mathbf{w})$ could be directly measured by asking the decision maker to contemplate small changes in her wealth in each state and to assess how her risk-neutral probabilities (i.e., her betting rates on individual states) would change as a result. If w_k is increased by a small amount Δw_k, the decision maker's risk-neutral probability in state j increases by $D\pi_{jk}(\mathbf{w})\Delta w_k + o(\Delta w_k)$, ceteris paribus, and when total wealth changes from $\mathbf{w}$ to $\mathbf{w} + \Delta\mathbf{w}$, her risk-neutral probability distribution changes from π to $\pi + \Delta\pi$, where $\Delta\pi = D\pi(\mathbf{w})\Delta\mathbf{w} + o(|\Delta\mathbf{w}|)$. $D\pi(\mathbf{w})$ is generally asymmetric and has less than full rank. In particular, its columns sum to zero, which guarantees that the solution to $\Delta\pi = D\pi(\mathbf{w})\Delta\mathbf{w}$ satisfies $\sum_{j=1}^{n} \Delta\pi_j = 0$, a necessary condition for no-arbitrage.

The main result of this section is that the decision maker's local and global attitudes toward risk and uncertainty are completely summarized by $D\pi(\mathbf{w})$, generalizing the state-independent expected-utility analysis of Pratt-Arrow and the two-dimensional state-preference analysis of Yaari:

PROPOSITION 2. (a) *The risk premium of a small neutral asset satisfies*

$$b(\mathbf{z};\mathbf{w}) \approx -\tfrac{1}{2}\mathbf{z} \cdot D\pi(\mathbf{w})\mathbf{z}.$$

(b) *The risk premium of a small nonneutral asset satisfies*

$$b(\mathbf{z};\mathbf{w}) \approx -\tfrac{1}{2}Q(\mathbf{z};\mathbf{w}), \quad \text{where}$$
$$Q(\mathbf{z};\mathbf{w}) \equiv (\mathbf{z} - \mathbf{z} \cdot \pi(\mathbf{w})) \cdot D\pi(\mathbf{w})(\mathbf{z} - \mathbf{z} \cdot \pi(\mathbf{w})).$$

(c) *The risk premium of any asset satisfies*

$$b(\mathbf{z};\mathbf{w}) = -\int_0^1 (1-x)Q(\mathbf{z};\mathbf{w} + x\mathbf{z} - B(x\mathbf{z};\mathbf{w}))\, dx.$$

Consequently,

(d) *The decision maker is risk averse if and only if, at every $\mathbf{w}$, $D\pi(\mathbf{w})$ is negative semidefinite in the subspace of neutral assets.*

A detailed proof is given in the Appendix, but an informal direct proof of parts (a) and (b) will be sketched here. If $\mathbf{z}$ is neutral ($\mathrm{E}_{\pi(\mathbf{w})}[\mathbf{z}] = 0$), its buying price is $-b(\mathbf{z};\mathbf{w})$ by definition. If the decision maker buys it at this price, thus keeping her utility constant, her final wealth will differ from her initial wealth by the vector amount $\mathbf{z} + b(\mathbf{z};\mathbf{w})$ and the first-order change in her risk-neutral probability distribution will be $D\pi(\mathbf{w})(\mathbf{z} + b(\mathbf{z};\mathbf{w}))$. Suppose that she purchases $\mathbf{z}$

in small, equal increments at the prevailing marginal prices, thus holding her utility roughly constant. In the process, her risk-neutral distribution will change linearly from $\boldsymbol{\pi}$ to $\boldsymbol{\pi} + D\boldsymbol{\pi}(\mathbf{w})(\mathbf{z} + b(\mathbf{z}; \mathbf{w}))$ and her marginal price of $\mathbf{z}$ will increase linearly from 0 to $\mathbf{z} \cdot D\boldsymbol{\pi}(\mathbf{w})(\mathbf{z} + b(\mathbf{z}; \mathbf{w}))$ as her wealth follows a locally quadratic trajectory along an indifference curve. The average marginal price is the midpoint, namely $\frac{1}{2}\mathbf{z} \cdot D\boldsymbol{\pi}(\mathbf{w})(\mathbf{z} + b(\mathbf{z}; \mathbf{w}))$, which is also the total price by the linearity of the price trajectory. Hence, the risk premium satisfies

$$-b(\mathbf{z}; \mathbf{w}) \approx \tfrac{1}{2}\mathbf{z} \cdot D\boldsymbol{\pi}(\mathbf{w})(\mathbf{z} + b(\mathbf{z}; \mathbf{w})).$$

If $\mathbf{z}$ is sufficiently small—in particular, if $|\mathbf{z} \cdot D\boldsymbol{\pi}(\mathbf{w})| \ll 1$—then the factor of $b(\mathbf{z}; \mathbf{w})$ inside the parentheses on the right-hand side is insignificant, whence $b(\mathbf{z}; \mathbf{w}) \approx -\frac{1}{2}\mathbf{z} \cdot D\boldsymbol{\pi}(\mathbf{w})\mathbf{z}$. If $\mathbf{z}$ is not neutral, the more general form $b(\mathbf{z}; \mathbf{w}) \approx -\frac{1}{2}Q(\mathbf{z}; \mathbf{w})$ of part (b) follows from the identity $b(\mathbf{z} + x; \mathbf{w}) = b(\mathbf{z}; \mathbf{w})$ with $x = -\mathbf{z} \cdot \boldsymbol{\pi}(\mathbf{w})$.

Thus, in the n-dimensional state-preference framework, the expression $r(x)\mathrm{E_p}[\mathbf{z}^2]$ in Pratt's risk premium formula for a neutral asset (1) is replaced by the more general matrix expression $-\mathbf{z} \cdot D\boldsymbol{\pi}(\mathbf{w})\mathbf{z}$. Evidently, $D\boldsymbol{\pi}(\mathbf{w})$ encodes both the decision maker's beliefs and local risk preferences, and indeed it can be factored into a product of two matrices, one of which contains the decision maker's risk-neutral probabilities and the other of which is constructed from ratios of second and first derivatives of the ordinal utility function, generalizing the Pratt-Arrow measure. To show this, define the local *risk aversion matrix* as the matrix $\mathbf{R}(\mathbf{w})$ whose jkth element is the following ratio of second to first derivatives:

$$r_{jk}(\mathbf{w}) = -(\partial^2 U(\mathbf{w})/\partial w_j w_k)/(\partial U(\mathbf{w})/\partial w_j).$$

Under expected-utility preferences, $\mathbf{R}(\mathbf{w})$ would be an observable diagonal matrix, and with state-independent utility and riskless prior wealth $\mathbf{w} = x\mathbf{1}$, the diagonal elements would be $r_{jj}(x\mathbf{1}) = r(x)$ for every j, as will be discussed in more detail in the following section. But under general preferences, $\mathbf{R}(\mathbf{w})$ is neither a diagonal matrix nor is it observable, because it is not invariant to monotonic transformations of U. In particular, if $\widehat{U}(\mathbf{w}) = f(U(\mathbf{w}))$, where f is monotonic and twice differentiable, then $\widehat{U}$ represents the

same preferences as U, but the corresponding risk aversion matrix $\widehat{\mathbf{R}}(\mathbf{w})$ differs from $\mathbf{R}(\mathbf{w})$ by an additive constant in each column:

$$\widehat{\mathbf{R}}(\mathbf{w}) = \mathbf{R}(\mathbf{w}) + \alpha\overline{\Pi}(\mathbf{w}),$$

where $\overline{\Pi}(\mathbf{w})$ is the matrix whose rows are all equal to $\boldsymbol{\pi}(\mathbf{w})^{\mathrm{T}}$, i.e., the matrix whose elements in the kth column are all equal to $\boldsymbol{\pi}_k(\mathbf{w})$; and $\alpha = (\mathbf{1} \cdot \nabla U(\mathbf{w})) \times (f''(U(\mathbf{w}))/f'(U(\mathbf{w})))$ where f' and f'' are the first and second derivatives of f. To eliminate the arbitrary constants, let a *normalized risk aversion matrix* $\overline{\mathbf{R}}(\mathbf{w})$ be defined by

$$\overline{\mathbf{R}}(\mathbf{w}) = \mathbf{R}(\mathbf{w}) - \overline{\Pi}(\mathbf{w})\mathbf{R}(\mathbf{w}).$$

The jkth element of $\overline{\mathbf{R}}(\mathbf{w})$ is then

$$\bar{r}_{jk}(\mathbf{w}) = r_{jk}(\mathbf{w}) - \mathrm{E}_{\boldsymbol{\pi}(\mathbf{w})}[\mathbf{r}_k(\mathbf{w})], \qquad (3)$$

where $\mathbf{r}_k(\mathbf{w})$ denotes the kth column of $\mathbf{R}(\mathbf{w})$. It follows that $\mathrm{E}_{\boldsymbol{\pi}(\mathbf{w})}[\overline{\mathbf{R}}(\mathbf{w})\Delta\mathbf{w}] = \mathbf{0}$ for any vector $\Delta\mathbf{w}$. The normalized risk aversion matrix is invariant to monotonic transformations of U and is observable. Comparison of terms in (2) and (3) reveals that $D\boldsymbol{\pi}(\mathbf{w})$ and $\overline{\mathbf{R}}(\mathbf{w})$ are related by

$$D\boldsymbol{\pi}(\mathbf{w}) = -\Pi(\mathbf{w})\overline{\mathbf{R}}(\mathbf{w}), \qquad (4)$$

where $\Pi(\mathbf{w}) = \mathrm{diag}(\boldsymbol{\pi}(\mathbf{w}))$. From (4) it is seen that the jkth element of $\overline{\mathbf{R}}(\mathbf{w})$ is $-(\partial\boldsymbol{\pi}_j(w)/\partial w_k)/\boldsymbol{\pi}_j(w)$, which is minus the *relative* rate of change of the risk-neutral probability of state j as wealth increases in state k. In these terms, we have

COROLLARY 2.1. *The risk premium of a small neutral asset* $\mathbf{z}$ *satisfies*

$$b(\mathbf{z}; \mathbf{w}) \approx \tfrac{1}{2}\mathbf{z} \cdot \Pi(\mathbf{w})\overline{\mathbf{R}}(\mathbf{w})\mathbf{z} = \tfrac{1}{2}\mathbf{z} \cdot \Pi(\mathbf{w})\mathbf{R}(\mathbf{w})\mathbf{z}$$
$$= \tfrac{1}{2}\mathrm{Cov}_{\boldsymbol{\pi}(\mathbf{w})}[\mathbf{z}, \mathbf{R}(\mathbf{w})\mathbf{z}].$$

PROOF. The first identity follows from (4). The second follows from the fact that $\mathbf{R}$ can be substituted for $\overline{\mathbf{R}}$ when $\mathbf{z}$ is neutral, as it differs only by the columnwise addition of constants that drop out when it is premultiplied by $\mathbf{z} \cdot \Pi(\mathbf{w})$. □

Comparison with (1) shows that, under general conditions, the "true" distribution $\mathbf{p}$ in the risk premium formula is replaced by the local risk-neutral distribution $\boldsymbol{\pi}(\mathbf{w})$, while the scalar Pratt-Arrow measure $r(x)$ is replaced by the matrix risk aversion measure $\mathbf{R}(\mathbf{w})$, or equivalently by its normalized, observable form $\overline{\mathbf{R}}(\mathbf{w})$.

4. The Special Case of Separable Preferences (State-Dependent Utility)

In the special case where the decision maker has preferences that are separable across mutually exclusive events—i.e., preferences that satisfy the independence axiom[8]—her utility function has the cardinal, additively separable representation:

$$U(\mathbf{w}) = v_1(w_1) + \cdots + v_n(w_n)$$

(Debreu 1960, Fishburn and Wakker 1995). This representation of preferences is equivalent to subjective expected utility with state-dependent utilities and not-necessarily-unique subjective probabilities, for it is always possible to write $v_j(w_j) = p_j u_j(w_j)$ where the numbers $\{p_j\}$ are arbitrarily chosen probabilities summing to 1 and the functions $\{u_j(w_j)\}$ are correspondingly scaled state-dependent utilities.[9] Because events are not uniquely ordered by probability under this representation, separability of preferences is not sufficient for probabilistic sophistication, which Epstein (1999, also Epstein and Zhang 2001) has equated with uncertainty neutrality in a Savage-act framework. However, in the present framework separability is sufficient for uncertainty neutrality because it ensures that preferences have an expected-utility representation even if it is not unique.

When U is additively separable, its cross-derivatives are zero and $\mathbf{R}(\mathbf{w}) = \mathrm{diag}(\mathbf{r}(\mathbf{w}))$, where $\mathbf{r}(\mathbf{w})$ is a vector-valued Pratt-Arrow measure of risk aversion whose jth element is

$$r_j(\mathbf{w}) = -(\partial^2 U(\mathbf{w})/\partial w_j^2)/(\partial U(\mathbf{w})/\partial w_j)$$
$$= -u_j''(w_j)/u_j'(w_j), \tag{5}$$

and u_j is the Bernoulli utility function for money in state j in an arbitrary expected-utility representation. (The confounded probabilities and utility-scale factors conveniently drop out when $\mathbf{r}$ is computed.) Correspondingly, the elements of $\overline{\mathbf{R}}(\mathbf{w})$ satisfy $\bar{r}_{jk}(\mathbf{w}) = r_k(\mathbf{w})(1_{jk} - \boldsymbol{\pi}_k(\mathbf{w}))$, which can be inverted to obtain $r_j(\mathbf{w}) = \bar{r}_{jj}(\mathbf{w}) - \bar{r}_{kj}(\mathbf{w})$ for $j \neq k$, whence $\mathbf{r}(\mathbf{w})$, like $\overline{\mathbf{R}}(\mathbf{w})$, is directly observable. The matrix $D\boldsymbol{\pi}(\mathbf{w})$ of derivatives of the risk-neutral probabilities has the generic element

$$D\boldsymbol{\pi}_{jk}(\mathbf{w}) = -\boldsymbol{\pi}_j(\mathbf{w})r_k(\mathbf{w})(1_{jk} - \boldsymbol{\pi}_k(\mathbf{w})),$$

and the risk premium approximation formula is accordingly specialized as:

COROLLARY 2.2. *For a decision maker with separable preferences, the risk premium of a small neutral asset* $\mathbf{z}$ *is*

$$b(\mathbf{z}; \mathbf{w}) \approx \tfrac{1}{2}\mathrm{E}_{\boldsymbol{\pi}(\mathbf{w})}[\mathbf{r}(\mathbf{w})\mathbf{z}^2] = \tfrac{1}{2}\mathrm{Cov}_{\boldsymbol{\pi}(\mathbf{w})}[\mathbf{z}, \mathbf{r}(\mathbf{w})\mathbf{z}].$$

It follows that a sufficient condition for the decision maker to be risk averse is $\mathbf{r}(\mathbf{w}) \geq 0$ at every $\mathbf{w}$, and a necessary condition is that *at most one* element of $\mathbf{r}(\mathbf{w})$ may be negative.[10]

By taking expectations with respect to the observable risk-neutral distribution rather than the unobservable true distribution, the problems of stochastic prior wealth and state-dependent utility have been finessed away: The decision maker's true probabilities and the correlations between the risky asset and her prior wealth are irrelevant once $\boldsymbol{\pi}(\mathbf{w})$ and $\mathbf{r}(\mathbf{w})$ have been observed. In the special case where the Pratt-Arrow measure is constant across states (e.g., if the decision maker has state-independent exponential

---

[8] Let $(\mathbf{x}_A, \mathbf{y}_{\sim A})$ denote the wealth distribution that agrees with $\mathbf{x}$ in event A (a subset of states) and agrees with $\mathbf{y}$ otherwise. The independence axiom, which is Savage's (1954) Postulate P2, requires that $(\mathbf{x}_A, \mathbf{y}_{\sim A}) \succeq (\mathbf{x}'_A, \mathbf{y}_{\sim A})$ if and only if $(\mathbf{x}_A, \mathbf{y}'_{\sim A}) \succeq (\mathbf{x}'_A, \mathbf{y}'_{\sim A})$ for all $\mathbf{x}$, $\mathbf{y}$, $\mathbf{x}'$, $\mathbf{y}'$, and every event A. In other words, if two wealth distributions agree on some subset of states, then the direction of preference between them doesn't depend on *how* they agree there. A behavioral violation of the independence axiom could be due to a dislike of ambiguous probabilities (as in Ellsberg's paradox) or some other cause (e.g., an attraction to sure things, as in Allais' paradox).

[9] Even under Savage's axioms, preferences among material acts do not uniquely determine subjective probabilities when utilities are potentially state dependent, a troublesome issue that has been discussed by Aumann (1971), Shafer (1986), Schervish et al. (1990), Karni and Mongin (2000), and Nau (1995, 2001) among others. Notwithstanding, Karni (1983, 1985) and Kelsey and Nordquist (1991) require that probabilities be uniquely determined for purposes of defining risk aversion under state-dependent utility, e.g., by the method of Karni et al. (1983).

[10] It is permissible for the Pratt-Arrow measure to be negative in one state because a gamble that yields a nonzero payoff in that state must also yield a nonzero payoff in one or more other states, and if the Pratt-Arrow measure in all other states is sufficiently positive, the net effect is still risk aversion.

utility), $\mathbf{R}(\mathbf{w}) \equiv \mathrm{diag}(r)$ where r is a scalar, in which case $D\boldsymbol{\pi}(\mathbf{w})$ is symmetric and has zero row sums as well as zero column sums. The risk aversion measure can then be taken outside the expectation in the approximation formula for the local risk premium:

$$b(\mathbf{z}; \mathbf{w}) \approx \tfrac{1}{2} r \mathrm{E}_{\boldsymbol{\pi}(\mathbf{w})}[\mathbf{z}^2] = \tfrac{1}{2} r \, \mathrm{Var}_{\boldsymbol{\pi}(\mathbf{w})}[\mathbf{z}].$$

However, even here the local risk-neutral distribution $\boldsymbol{\pi}(\mathbf{w})$, rather than the true distribution $\mathbf{p}$, is used to evaluate the variance when marginal utilities vary across states due to prior stakes. The exact risk premium similarly satisfies

$$b(\mathbf{z}; \mathbf{w}) = r \int_0^1 (1-x) \, \mathrm{Var}_{\boldsymbol{\pi}(\mathbf{w}+x\mathbf{z}-B(x\mathbf{z};\,\mathbf{w}))}[\mathbf{z}] \, dx,$$

which is a weighted average of $\tfrac{1}{2} r \, \mathrm{Var}_{\boldsymbol{\pi}(\mathbf{w})}[\mathbf{z}]$ along the indifference curve through $\mathbf{w}$, so that the risk premium is transparently the price of variance ($\tfrac{1}{2} r$) multiplied by a weighted average of the *risk-neutral* variance of $\mathbf{z}$ in the vicinity of current wealth.

The results of this section can be summarized as follows: If the decision maker has convex preferences that are separable across states, her local preferences under uncertainty are completely characterized by (at most) a pair of numbers for every state—a risk-neutral probability and a risk aversion coefficient. Such a person is "risk averse but uncertainty neutral." The risk aversion coefficients may be state dependent, but—unlike subjective probabilities or utilities—they are uniquely determined by preferences, and hence they are observable. If separability does not hold, the more general matrix representation of Proposition 2 applies, and (like Dan in Figure 1) the decision maker may exhibit aversion to uncertainty as well as risk.

5. Marginal Investment and Comparative Risk Aversion

Under state-independent expected-utility preferences, it is possible to describe comparative risk aversion and comparative investment behavior in terms of properties of the Bernoulli utility function, as summarized by its Pratt-Arrow measure. Naturally, Agent 1 is "more risk averse" than Agent 2 if $r_1(x) > r_2(x)$ at every level of riskless wealth x, where r_i is the Pratt-Arrow measure of agent i. If Agent 1 is

more risk averse than Agent 2 in this sense, she will purchase smaller quantities of any risky asset than Agent 2 when both start from the same riskless initial wealth position. When prior wealth is risky (stochastic), the situation is more complicated, and stronger notions of comparative risk aversion and restrictions on the joint distributions of old and new risks are needed to obtain similar results. The most tractable and thoroughly studied case is that of probabilistic independence between prior wealth and new risks (e.g., Kihlstrom et al. 1981, Pratt 1988, Gollier and Pratt 1996, Eeckhoudt et al. 1996; a different approach is taken by Ross 1981). Significantly, when two expected-utility-maximizing decision makers with the same probabilistic beliefs contemplate the same risky asset that is independent of their prior wealth, they will agree on the risk-neutral distribution of the new asset: Their risk-neutral distributions for the asset will simply coincide with its assumed true distribution because their expected marginal utilities for money will not depend on the new asset's value.

In the much more general setting considered in this paper, it is not possible to surgically remove a Bernoulli utility function from the decision maker: Risk premia and investment behavior may depend on unobserved stochasticity of prior wealth, state-dependence of utility, and aversion to uncertainty as well as risk. Nevertheless, in the spirit of Yaari's (1969) characterization of "more risk averse," the local and nonlocal risk preferences of different decision makers can be compared in terms of the curvature of their indifference curves for wealth, as quantified by the matrices of derivatives of their risk-neutral probabilities in appropriate neighborhoods of the status quo. It is not necessary for decision makers to agree on the probability distributions of assets—or even to have probabilistic beliefs—although when comparing risk attitudes toward a particular asset it will be necessary to assume that they initially agree on the asset's marginal price—i.e., its risk-neutral expectation. They can always reach such an agreement, if necessary, by trading the asset between them, even if other risks that they face are not insurable.

The local curvature of indifference curves in the direction of $\mathbf{z}$ is measured by the quadratic form $\mathbf{z} \cdot D\boldsymbol{\pi}(\mathbf{w})\mathbf{z}$. Optimal purchases of a single risky asset are

inversely proportional to this measure of curvature, appropriately averaged, as shown in the following:

PROPOSITION 3. *If a risk-averse individual has the opportunity (only) to purchase shares of a new asset with net payoff vector* $\mathbf{z}$ *having a positive marginal price* ($P(\mathbf{z}; \mathbf{w}) = \mathbf{z} \cdot \boldsymbol{\pi}(\mathbf{w}) > 0$):

(a) *She will optimally purchase a quantity* α *that satisfies*

$$- \int_0^\alpha \mathbf{z} \cdot D\boldsymbol{\pi}(\mathbf{w} + x\mathbf{z})\mathbf{z} \, dx = \mathbf{z} \cdot \boldsymbol{\pi}(\mathbf{w})$$

and

(b) *if* $(\mathbf{z} \cdot \boldsymbol{\pi}(\mathbf{w}))/\|\mathbf{z}\|$ *is sufficiently small, the optimal quantity satisfies* $\alpha \approx -(\mathbf{z} \cdot \boldsymbol{\pi}(\mathbf{w}))/(\mathbf{z} \cdot D\boldsymbol{\pi}(\mathbf{w})\mathbf{z})$.

PROOF. For part (a), note that the marginal price of $\mathbf{z}$ must be zero following an optimal purchase. As wealth changes from $\mathbf{w} + x\mathbf{z}$ to $\mathbf{w} + (x + dx)\mathbf{z}$, the risk-neutral distribution changes by $D\boldsymbol{\pi}(\mathbf{w} + x\mathbf{z})\mathbf{z} \, dx$, and the marginal price of $\mathbf{z}$ changes by $\mathbf{z} \cdot D\boldsymbol{\pi}(\mathbf{w} + x\mathbf{z})\mathbf{z} \, dx$. Hence, the total change in marginal price when $\alpha\mathbf{z}$ is acquired is $\int_0^\alpha \mathbf{z} \cdot D\boldsymbol{\pi}(\mathbf{w} + x\mathbf{z})\mathbf{z} \, dx$, and at the optimal value of α, this quantity must equal $-\mathbf{z} \cdot \boldsymbol{\pi}(\mathbf{w})$. The approximation formula in part (b) applies when the relative marginal price $(\mathbf{z} \cdot \boldsymbol{\pi}(\mathbf{w}))/\|\mathbf{z}\|$, which is the sine of the angle between $\mathbf{z}$ and its projection in the subspace of neutral assets, is small enough for $D\boldsymbol{\pi}(\mathbf{w})$ to be effectively constant over the range of integration.

COROLLARY 3.1. *Suppose that two individuals with wealth, prices, risk premia, etc., subscripted by* $i \in \{1, 2\}$, *have the opportunity to purchases shares of a new asset with net payoff vector* $\mathbf{z}$ *such that* $P_1(\mathbf{z}; \mathbf{w}_1) = P_2(\mathbf{z}; \mathbf{w}_2) > 0$, *and Individual* 1 *optimally purchases quantity* α. *Then Individual* 2 *will optimally purchase no less than individual* 1 *if* $\mathbf{z} \cdot D\boldsymbol{\pi}_1(\mathbf{w}_1 + x\mathbf{z})\mathbf{z} \leq \mathbf{z} \cdot D\boldsymbol{\pi}_2(\mathbf{w}_2 + x\mathbf{z})\mathbf{z}$[11] *for all* $x \leq \alpha$.

COROLLARY 3.2. *If* $\boldsymbol{\pi}_1(\mathbf{w}_1) = \boldsymbol{\pi}_2(\mathbf{w}_2)$ *and* $D\boldsymbol{\pi}_1(\mathbf{w}_1) - D\boldsymbol{\pi}_2(\mathbf{w}_2)$ *is negative semidefinite, then Individual* 2 *will purchase no less than Individual* 1 *of any asset for which*

---

[11] These quadratic forms are typically negative for risk-averse individuals, so the inequality in Corollary 3.1 means that the left-hand side is more negative (and hence larger in magnitude) than the right-hand side. The sense of both corollaries is that Individual 1 is more risk averse than Individual 2 if $D\boldsymbol{\pi}_1$ is "more negative semidefinite" than $D\boldsymbol{\pi}_2$ in an appropriate neighborhood of current wealth.

$(\mathbf{z} \cdot \boldsymbol{\pi}_i(\mathbf{w}_i))/\|\mathbf{z}\|$ *is sufficiently small. If both agents have separable preferences, a sufficient condition for the negative semidefiniteness requirement is* $\mathbf{r}_1(\mathbf{w}_1) \geq \mathbf{r}_2(\mathbf{w}_2)$ *pointwise, where* $\mathbf{r}_i$ *is the vector risk aversion measure of agent i.*

6. Markets under Uncertainty

If the decision maker is embedded in a market for contingent claims where the relative prices of assets are summarized by a market risk-neutral distribution $\boldsymbol{\pi}^*$, the first-order condition of utility maximization is that her own risk-neutral probabilities should equilibrate with those of the market; i.e., $\boldsymbol{\pi}(\mathbf{w}) = \boldsymbol{\pi}^*$.[12] Under these conditions, her responses to small changes in wealth or prices are completely determined by $\boldsymbol{\pi}(\mathbf{w})$ and $D\boldsymbol{\pi}(\mathbf{w})$. For example, suppose the decision maker receives a lump-sum amount of income Δx. It is of interest to determine the change $\Delta\mathbf{w}$ in her state-contingent wealth that will obtain after she has re-equilibrated with the market, holding prices fixed. This vector lies along the *wealth expansion path* emanating from the decision maker's current wealth distribution, and it is the solution of the equations $D\boldsymbol{\pi}(\mathbf{w})\Delta\mathbf{w} = 0$ (risk-neutral probabilities remain unchanged) and $\boldsymbol{\pi}(\mathbf{w}) \cdot \Delta\mathbf{w} = \Delta x$ (the market value of the wealth change is Δx). Another quantity of interest is the self-financed change in wealth $\Delta\mathbf{w}$ that will be observed if the market's risk-neutral probabilities change by a small amount $\Delta\boldsymbol{\pi}$, analogous to the Slutsky equation of consumer theory. The change in wealth is the solution to $D\boldsymbol{\pi}(\mathbf{w} + \frac{1}{2}\Delta\mathbf{w})\Delta\mathbf{w} = \Delta\boldsymbol{\pi}$ (risk-neutral probabilities change by $\Delta\boldsymbol{\pi}$, using the midpoint value of $D\boldsymbol{\pi}$ between $\mathbf{w}$ and $\mathbf{w} + \Delta\mathbf{w}$) and $(\boldsymbol{\pi}(\mathbf{w}) + \Delta\boldsymbol{\pi}) \cdot \Delta\mathbf{w} = 0$ (the new market value of the wealth change is zero). These details of these solutions are given by

PROPOSITION 4. *In a complete market for contingent claims*:

(a) *the change in wealth induced by a small lump-sum income* Δx, *holding prices fixed, is*

$$\Delta\mathbf{w} \approx \frac{\mathbf{R}^{-1}(\mathbf{w})\mathbf{1}}{\mathrm{E}_{\boldsymbol{\pi}(\mathbf{w})}[\mathbf{R}^{-1}(\mathbf{w})\mathbf{1}]}\Delta x,$$

---

[12] Schlee and Schlesinger's (1993) concept of a "generalized" risk premium is defined in terms of the value to the decision maker of the opportunity to trade contingent claims at exogenous market prices.

hence, $\mathbf{R}^{-1}(\mathbf{w})\mathbf{1}$ *is the direction of the wealth expansion path at* $\mathbf{w}$; *and*

(b) *the self-financed change in wealth induced by a small change* $\Delta\boldsymbol{\pi}$ *in market risk-neutral probabilities is*

$$\Delta\mathbf{w} \approx \mathbf{R}^{-1}(\mathbf{w})\left(\frac{-\Delta\boldsymbol{\pi}}{\boldsymbol{\pi}(\mathbf{w})+\frac{1}{2}\Delta\boldsymbol{\pi}}+y\mathbf{1}\right),$$

where

$$y \equiv \frac{\mathrm{E}_{\boldsymbol{\pi}(\mathbf{w})+\Delta\boldsymbol{\pi}}\left[\mathbf{R}^{-1}(\mathbf{w})\dfrac{\Delta\boldsymbol{\pi}}{\boldsymbol{\pi}(\mathbf{w})+\frac{1}{2}\Delta\boldsymbol{\pi}}\right]}{\mathrm{E}_{\boldsymbol{\pi}(\mathbf{w})+\Delta\boldsymbol{\pi}}[\mathbf{R}^{-1}(\mathbf{w})\mathbf{1}]}.$$

PROOF. For part (a), recall that $D\boldsymbol{\pi}(\mathbf{w}) = -\Pi(\mathbf{w}) \times (\mathbf{R}(\mathbf{w})-\overline{\Pi}(\mathbf{w})\mathbf{R}(\mathbf{w}))$ and verify by substitution that the wealth expansion condition, $D\boldsymbol{\pi}(\mathbf{w})\Delta\mathbf{w}=0$, is satisfied by $\Delta\mathbf{w} \propto \mathbf{R}^{-1}(\mathbf{w})\mathbf{1}$. The market value condition, $\boldsymbol{\pi}(\mathbf{w})\cdot\Delta\mathbf{w}=\Delta x$, is then satisfied by choosing $\Delta x/\mathrm{E}_{\boldsymbol{\pi}(\mathbf{w})}[\mathbf{R}^{-1}(\mathbf{w})\mathbf{1}]$ as the scale factor. For part (b), note that the term involving y is in the direction of the wealth expansion path (by part (a)), so it has no effect on risk-neutral probabilities, and y is determined precisely so that the self-financing condition, $(\boldsymbol{\pi}(\mathbf{w})+\Delta\boldsymbol{\pi})\cdot\Delta\mathbf{w}=0$, is satisfied. It remains to show that the term involving $-\Delta\boldsymbol{\pi}/(\boldsymbol{\pi}(\mathbf{w})+\frac{1}{2}\Delta\boldsymbol{\pi})$ yields the desired change $\Delta\boldsymbol{\pi}$ in risk-neutral probabilities, i.e., that $D\boldsymbol{\pi}(\mathbf{w}+\frac{1}{2}\Delta\mathbf{w})\mathbf{R}^{-1}(\mathbf{w})(-\Delta\boldsymbol{\pi}/(\boldsymbol{\pi}(\mathbf{w})+\frac{1}{2}\Delta\boldsymbol{\pi})) \approx \Delta\boldsymbol{\pi}$. This, in turn, can be verified through substitution by expressing the midpoint value of the derivative matrix, $D\boldsymbol{\pi}(\mathbf{w}+\frac{1}{2}\Delta\mathbf{w})$, as $-(\Pi(\mathbf{w})+\frac{1}{2}\Delta\Pi(\mathbf{w}))(\mathbf{R}(\mathbf{w})-(\overline{\Pi}(\mathbf{w})+\frac{1}{2}\Delta\overline{\Pi}(\mathbf{w}))\mathbf{R}(\mathbf{w}))$, on the assumption that $\mathbf{R}(\mathbf{w})$ is effectively constant (which is equivalent to ignoring third-order effects). □

The matrix $\mathbf{R}^{-1}(\mathbf{w})$, which measures the decision maker's local *risk tolerance*, plays a prominent role in these results.[13] Part (a) establishes that $\mathbf{R}^{-1}(\mathbf{w})$ is the linear transformation that maps small changes in income into changes in the optimal distribution of wealth; thus, the decision maker prefers to redistribute income across states in proportion to risk tolerance. Part (b) establishes that $\mathbf{R}^{-1}(\mathbf{w})$ is also the linear

transformation that maps small changes in relative market prices into changes in the optimal distribution of wealth. It is suggestive to rewrite the formula in part (b) as follows:

$$\mathbf{R}(\mathbf{w})\Delta\mathbf{w} \approx \frac{-\Delta\boldsymbol{\pi}}{\boldsymbol{\pi}(\mathbf{w})+\frac{1}{2}\Delta\boldsymbol{\pi}}+y. \qquad (6)$$

The term on the left can be interpreted as the change in wealth expressed in dimensionless, risk-adjusted units, i.e., units of "wealth divided by risk tolerance." By virtue of the result in part (a), the constant term y on the right has no effect on the decision maker's risk-neutral probabilities, as it corresponds to a move along the wealth expansion path (assuming $\mathbf{R}(\mathbf{w}) \approx \mathbf{R}(\mathbf{w}+\Delta\mathbf{w})$). Hence, the term $-\Delta\boldsymbol{\pi}/(\boldsymbol{\pi}(\mathbf{w})+\frac{1}{2}\Delta\boldsymbol{\pi})$, which is the negative of the relative change in prices, is entirely responsible for producing the required change in the decision maker's risk-neutral probabilities. Now, intuitively, the decision maker should react to movements in prices by shifting wealth away from states where the relative prices of contingent claims have increased. The rewritten formula (6) shows that equilibrium with the market is restored by making a change in risk-adjusted wealth that is exactly *equal and opposite to the relative change in prices*. The constant term y merely recenters the transaction so that it is self-financing at the new prices.

When preferences are separable, the preceding results can be simplified by replacing the risk aversion matrix $\mathbf{R}(\mathbf{w})$ with the (observable) risk aversion vector $\mathbf{r}(\mathbf{w})$. Correspondingly, the risk tolerance matrix $\mathbf{R}^{-1}(\mathbf{w})$ is replaced by the *risk tolerance vector* $\mathbf{t}(\mathbf{w}) \equiv \mathbf{r}(\mathbf{w})^{-1}$.

COROLLARY 4.1. *For a decision maker with separable preferences in a complete market for contingent claims:*

(a) *the change in wealth induced by a small lump-sum income* Δx, *holding prices fixed, is*

$$\Delta\mathbf{w} \approx \frac{\mathbf{t}(\mathbf{w})}{\mathrm{E}_{\boldsymbol{\pi}(\mathbf{w})}[\mathbf{t}(\mathbf{w})]}\Delta x,$$

hence the risk tolerance vector $\mathbf{t}(\mathbf{w})$ *is the direction of the wealth expansion path at* $\mathbf{w}$; *and*

(b) *the self-financed change in wealth induced by a small change* $\Delta\boldsymbol{\pi}$ *in market risk-neutral probabilities is*

$$\Delta\mathbf{w} \approx \mathbf{t}(\mathbf{w})\left(\frac{-\Delta\boldsymbol{\pi}}{\boldsymbol{\pi}(\mathbf{w})+\frac{1}{2}\Delta\boldsymbol{\pi}}+y\right),$$

---

[13] The formulas in the proposition are stated in terms of the nonsingular but unobservable matrix $\mathbf{R}$, rather than its observable but singular counterpart $\overline{\mathbf{R}} = \mathbf{R}-\overline{\Pi}\mathbf{R}$. However, the formulas are invariant to the addition of constants to columns of $\mathbf{R}$, so $\mathbf{R}$ could be replaced by $\overline{\mathbf{R}}+\mathbf{A}$, where $\mathbf{A}$ is any matrix with nonzero constant columns.

where

$$y \equiv \frac{E_{\pi(\mathbf{w})+\Delta\pi}\left[\mathbf{t}(\mathbf{w})\frac{\Delta\pi}{\pi(\mathbf{w})+\frac{1}{2}\Delta\pi}\right]}{E_{\pi(\mathbf{w})+\Delta\pi}[\mathbf{t}(\mathbf{w})]}.$$

7. An Application to Project Valuation

Proposition 4 and its corollary are potentially applicable to the problem of choosing among alternative risky projects in the setting of a complete market for contingent claims. A project is characterized by a stream of state- and time-dependent cash flows. The decision maker (e.g., a firm) is assumed to have a utility function whose arguments are state- and time-dependent amounts of consumption, and asset purchases can be used to transfer consumption across time and states. If the market is complete, the solution to the project selection problem does not depend on the decision maker's own risk preferences. Rather, the optimal project is the one whose cash flow stream has the highest expected net present value, where the market risk-free interest rate is used for purposes of discounting and the market risk-neutral probabilities are used for computing expected values (Ross 1976, 1978; Rubinstein 1976). The decision maker's risk preferences play a role only in the financing problem—i.e., given the optimal project, what additional asset purchases should be made for optimal borrowing and risk hedging? The financing problem is solved by purchasing assets so as to restore the equilibrium between the decision maker's risk-neutral probabilities and those of the market, after the project has been added to the existing portfolio. If the decision maker is in equilibrium with the market prior to the project decision, the optimal financing decision is to "short" the project cash flows by selling a replicating portfolio and then invest the arbitrage profit according to the formula in part (a) of the proposition or its corollary. If the decision maker is not yet in equilibrium with the market, the optimal financing decision can be (approximately) determined by applying the formula in part (b) after adding the project to current wealth.

The same modeling framework applies equally well to intertemporal decision problems, where wealth varies across time as well as states, and the elements of the vector $\boldsymbol{\pi}(\mathbf{w})$ are more appropriately called *normalized state prices* rather than risk-neutral probabilities. As an example of an intertemporal application, consider the simple capital budgeting problem due to Trigeorgis and Mason (1987) and discussed by Nau and McCardle (1991) and Smith and Nau (1995). There are two dates (0 and 1) and two states of the world ("good" and "bad" returns on investment) at Date 1. The firm has three decision alternatives with respect to the construction of a new plant: "invest," "defer," and "decline." The firm can also buy or sell two securities, one risky and one risk-free. The net cash flow streams associated with the alternatives and the securities are as presented in Table 1.

Let a consumption stream be represented by a corresponding three-vector (w_0, w_{1g}, w_{1b}). The firm's utility function is assumed to have the time-additive exponential form

$$U(w_0, w_{1g}, w_{1b})$$
$$= -\exp(-w_0/200) - 0.5\exp(-w_{1g}/220)$$
$$- 0.5\exp(-w_{1b}/220).$$

Thus, it is as if the firm assigns equal probability to the two states at Date 1 and has constant absolute risk aversion with risk tolerances of 200 and 220 at Dates 0 and 1, respectively. From the asset prices and the utility function, assuming constant prior wealth, the state prices can be derived for the market and for the firm, the latter depending on the alternative chosen (see Table 2).

Notice that the firm's state price distribution (the normalized gradient of $U(\mathbf{w})$) is fairly close to that of the market under the "Defer" and "Decline" alternatives (absent any asset purchases), but it deviates sharply under the "Invest" alternative. The optimal

Table 1   Cash Flows of Projects and Unit Asset Purchases

	Date 0	Date 1, good state	Date 1, bad state
Invest	−104	180	60
Defer	0	67.68	0
Decline	0	0	0
Buy 1 share of risk-free asset	−1	1.08	1.08
Buy 1 share of risky asset	−20	36	12

Table 2   Normalized State Prices for Market and Firm

	Date 0	Date 1, good state	Date 1, bad state
Market	0.519	0.192	0.289
Firm w/invest	0.755	0.090	0.155
Firm w/defer	0.559	0.187	0.254
Firm w/decline	0.524	0.238	0.238

overall strategy for the firm is to choose the alternative whose cash flow stream has the highest value under the market's state prices and then purchase assets so as to equalize the firm's state prices with those of the market. Table 3 compares the market values of the three alternatives, as well as the exactly and approximately optimal asset purchases under each alternative, where the approximation formula from Corollary 4.1(b) has been used. The "Defer" alternative is the optimal choice, and the approximation formula yields asset quantities very close to the exact values.

The approximation is slightly less good for the "Invest" alternative than for the others, because the necessary adjustment to state prices is much larger in that case, but nevertheless the approximation is so close that the buying risk premium for the difference in wealth distributions, which measures the impact of the error in monetary terms, is only 0.026—negligible in comparison to the project value.

If the firm is already in equilibrium with the market—i.e., in possession of the optimal asset position for the "Decline" alternative—part (a) of Corollary 4.1 can be applied instead to determine the change in the asset position that is needed when

Table 3   Project Valuations and Optimal Asset Purchases

	Market value		Optimal shares of Risk-free asset		Optimal shares of Risky asset	
	Normalized	Unnormalized	Exact	Approx.	Exact	Approx.
Invest	−2.077	−4.000	−80.24	−80.09	−1.283	−1.142
Defer	13.015	25.067	−34.24	−34.19	0.897	0.895
Decline	0	0	−78.22	−77.96	3.717	3.704

Note. The unnormalized value (which is the expected net present value at market probabilities and discount rates) is obtained by scaling the state prices so that the Date 0 state price is 1.

choosing the "Defer" option. The solution is to sell the "Defer" alternative at its market value (i.e., sell a replicating portfolio), then invest the income ($\Delta x = 13.015$ in normalized units) according to the formula in 4.1(a), i.e., redistribute the income across the three time-state contingencies in proportion to local risk tolerances (200, 220, and 220, respectively). Under constant absolute risk aversion, the latter method yields an exact solution to the financing problem, whereas the formula in 4.1(b) yields only an approximation (albeit a very good one). The point of this exercise is to show that, for a decision maker with separable preferences in a complete market, risk-neutral probabilities and statewise risk tolerances are sufficient statistics for choosing among small or moderate risks and determining optimal hedging strategies. Under more general preferences, the risk-neutral probabilities and their matrix of derivatives would be sufficient.

8.   Concluding Comments

The concepts of probabilistic beliefs, riskless wealth positions, and consequences with state-independent utility have traditionally played key roles in models of choice under uncertainty. However, a growing body of literature casts doubt on their uniqueness and observability, if not their existence, and consequently it is of interest to determine whether economic phenomena such as aversion to risk and uncertainty can be modeled without reference to them. Yaari's (1969) payoff-convexity definition of risk aversion, which does not refer to probabilities or riskless wealth, suggests that a simple and general measure of risk aversion ought to be available for a broad class of preferences under uncertainty. This paper has derived such a measure in terms of the matrix of derivatives of the decision maker's local risk-neutral probabilities. The measure applies to fairly general preferences, including expected-utility preferences that need not be state-independent and smooth nonexpected-utility preferences that need not be probabilistically sophisticated. It has been shown that various aspects of risk-averse behavior and financial decision making can be modeled entirely in terms of the decision

maker's risk-neutral probabilities, consistent with the central role that risk-neutral probabilities play elsewhere in models of markets under uncertainty (e.g., asset pricing by arbitrage). If the decision maker has separable preferences with stochastic prior wealth and/or state-dependent utility, her risk-neutral distribution is the appropriate distribution with which to compute the variance in Pratt's risk premium formula, and her local decision making behavior is completely characterized by the risk-neutral distribution and a vector measure of risk aversion. In these terms, risk aversion may be compared between individuals without observing or agreeing on true probabilities and without reference to the 45-degree certainty line, whose location may be unknown. Under more general smooth nonexpected utility preferences, the risk aversion measure is matrix-valued. A companion paper (Nau 2002) shows how aversion to uncertainty as well as risk may be encoded in the matrix of derivatives of the decision maker's risk-neutral probabilities.

Acknowledgments

Appendix

PROOF OF PROPOSITION 2. For part (c), consider the explicit sequence of wealth positions $\{\mathbf{w}_0, \mathbf{w}_1, \ldots, \mathbf{w}_m\}$ defined by $\mathbf{w}_0 = \mathbf{w}$ and $\mathbf{w}_i = \mathbf{w} + (i/m)\mathbf{z} - B((i/m)\mathbf{z}; \mathbf{w})$, which lie along the indifference curve from $\mathbf{w}$ to $\mathbf{w} + \mathbf{z} - B(\mathbf{z}; \mathbf{w})$ and are generated by the purchases of m equal increments of $\mathbf{z}$ at their respective buying prices, thus holding utility constant. Let $\boldsymbol{\pi}_i$ and $D\boldsymbol{\pi}_i$ denote the risk-neutral probabilities and their derivatives at $\mathbf{w}_i$, and let $\Delta \mathbf{w}_i = \mathbf{w}_i - \mathbf{w}_{i-1}$. Define a sequence of approximations to $\boldsymbol{\pi}_i$ by $\hat{\boldsymbol{\pi}}_0 = \boldsymbol{\pi}$ and

$$\hat{\boldsymbol{\pi}}_i = \hat{\boldsymbol{\pi}}_{i-1} + D\boldsymbol{\pi}_{i-1}\Delta\mathbf{w}_i = \boldsymbol{\pi} + \sum_{k=1}^{i-1} D\boldsymbol{\pi}_{k-1}\Delta\mathbf{w}_k$$

and note that (assuming continuity of $D\boldsymbol{\pi}$)

$$\boldsymbol{\pi}_i = \hat{\boldsymbol{\pi}}_i + o(m^{-1})$$

and

$$\Delta\mathbf{w}_i = \mathbf{z}/m - (\mathbf{z}/m)\cdot\boldsymbol{\pi}_{i-1} + o(m^{-1})$$
$$= \mathbf{z}/m - (\mathbf{z}/m)\cdot\hat{\boldsymbol{\pi}}_{i-1} + o(m^{-1}),$$

whence

$$b(\mathbf{z}; \mathbf{w}) = -B(\mathbf{z}; \mathbf{w}) = \sum_{i=1}^{m} \Delta\mathbf{w}_i - \mathbf{z}$$

$$= -\sum_{i=1}^{m-1}(m-i)(\mathbf{z}/m)\cdot D\boldsymbol{\pi}_{i-1}\Delta\mathbf{w}_i + O(m^{-1})$$

$$= -\sum_{i=1}^{m-1}(m-i)(\mathbf{z}/m)\cdot D\boldsymbol{\pi}_{i-1}(\mathbf{z}/m - (\mathbf{z}/m)\cdot\boldsymbol{\pi}_{i-1})$$

$$+ O(m^{-1}).$$

Next, $(\mathbf{z}/m)$ can be replaced by $\mathbf{z}/m - (\mathbf{z}/m)\cdot\boldsymbol{\pi}_{i-1}$ on the left side of $D\boldsymbol{\pi}_{i-1}$, because it differs only by a constant that is irrelevant because the column sums of $D\boldsymbol{\pi}_{i-1}$ are zero. This yields the quadratic form

$$(\mathbf{z}/m - (\mathbf{z}/m)\cdot\boldsymbol{\pi}_{i-1})\cdot D\boldsymbol{\pi}_{i-1}(\mathbf{z}/m - (\mathbf{z}/m)\cdot\boldsymbol{\pi}_{i-1})$$

$$= (\mathbf{z} - \mathbf{z}\cdot\boldsymbol{\pi}_{i-1})\cdot D\boldsymbol{\pi}_{i-1}(\mathbf{z} - \mathbf{z}\cdot\boldsymbol{\pi}_{i-1})/m^2,$$

which in turn is equal to to $Q(\mathbf{z}; \mathbf{w}_{i-1})/m^2$ by definition, whence

$$b(\mathbf{z}; \mathbf{w}) = -\sum_{i=1}^{m-1}(m-i)Q(\mathbf{z}; \mathbf{w}_{i-1})/m^2 + O(m^{-1})$$

$$= -\sum_{i=1}^{m-1}(1 - i/m)Q(\mathbf{z}; \mathbf{w}_{i-1})/m + O(m^{-1})$$

$$\rightarrow -\int_0^1 (1-x)Q(\mathbf{z}; \mathbf{w} + x\mathbf{z} - B(x\mathbf{z}; w))\, dx \quad \text{as } m \rightarrow \infty.$$

Parts (a) and (b) follow from part (c) by taking Q to be constant over the range of integration. The "if" part of (d) follows from (c), and the "only if" part follows from (a). □

References

Arrow, K. J. 1964. The role of securities in the optimal allocation of risk-bearing. *Quart. J. Econom.* **31** 91–96.

———. 1965. The theory of risk aversion. Lecture 2 in *Aspects of the Theory of Risk-Bearing*, Yrjo Jahnsson Lectures. Yrjo Jahnssonin Saatio, Helsinki, Finland.

Aumann, R. 1971. Letter to L. J. Savage. Reprinted in J. Drèze. 1987. *Essays on Economic Decision under Uncertainty*. Cambridge University Press, London, U.K.

Chew, S. H. 1995. A Schur-concave characterization of risk aversion for non-expected utility preferences. *J. Econom. Theory* **67** 402–435.

———, E. Karni, Z. Safra. 1987. Risk aversion in the theory of expected utility with rank dependent probabilities. *J. Econom. Theory* **42** 370–381.

Cohen, M. 1995. Risk aversion concepts in expected and non-expected utility theory. *Geneva Papers Risk Insurance Theory* **20**(1) 73–92.

Courtault, J.-M., J.-P. Gayant. 1998. Local risk aversion in the rank dependent expected utility model: First order versus second order effects. *Econom. Lett.* **59** 207–212.

de Finetti, B. 1952. Sulla preferibilitá. *Giornale Econom. Ann. Econom.* **11** 685–709.

Debreu, G. 1959. *Theory of Value*. Cowles Foundation, New Haven, CT.

——. 1960. Topological methods in cardinal utility theory. K. J. Arrow, S. Karlin, P. Suppes, eds. *Mathematical Methods in the Social Sciences*. Stanford University Press, Stanford, CA, 16–26.

Drèze, J. 1970. Market allocation under uncertainty. *Eur. Econom. Rev.* **2** 133–165.

Eeckhoudt, L., C. Gollier, H. Schlesinger. 1996. Changes in background risk and risk taking behavior. *Econometrica* **64** 683–689.

Epstein, L. 1999. A definition of uncertainty aversion. *Rev. Econom. Stud.* **66** 579–608.

——, J. Zhang. 2001. Subjective probabilities on subjectively unambiguous events. *Econometrica* **69** 265–306.

Fishburn, P., P. Wakker. 1995. The invention of the independence condition for preferences. *Management Sci.* **41** 1130–1144.

Gollier, C., J. W. Pratt. 1996. Risk vulnerability and the tempering effect of background risk. *Econometrica* **64** 1109–1123.

Hirshleifer, J. 1965. Investment decision under uncertainty: Choice-theoretic approaches. *Quart. J. Econom.* **74** 509–536.

Karni, E. 1983. Risk aversion for state-dependent utility functions: Measurement and applications. *Internat. Econom. Rev.* **24** 637–647.

——. 1985. *Decision Making under Uncertainty: The Case of State-Dependent Preferences*. Harvard University Press, Cambridge, MA.

——. 1995. Non-expected utility and the robustness of the classical insurance paradigm: Discussion. *Geneva Papers Risk Insurance Theory* **20** 51–56.

——, P. Mongin. 2000. On the determination of subjective probability by choices. *Management Sci.* **46** 233–248.

——, D. Schmeidler, K. Vind. 1983. On state dependent preferences and subjective probabilities. *Econometrica* **51** 1021–1031.

Kelsey, D., G. Nordquist. 1991. A more general measure of risk aversion when utility is state-dependent. *Oxford Econom. Papers* **43** 59–74.

Kihlstrom, R., D. Romer, S. Williams. 1981. Risk aversion with random initial wealth. *Econometrica* **49** 911–920.

Kimball, M. 1990. Precautionary saving in the small and in the large. *Econometrica* **58** 53–74.

Klibanoff, P., M. Marinacci, S. Mukerji. 2002. A smooth model of decision making under ambiguity. Working paper, Northwestern University, Evanston, IL.

Machina, M. 1995. Non-expected utility and the robustness of the classical insurance paradigm. *Geneva Papers Risk Insurance Theory* **20** 9–50.

——, D. Schmeidler. 1992. A more robust definition of subjective probability. *Econometrica* **60** 745–780.

Mas-Colell, A., M. Whinston, J. Green. 1995. *Microeconomic Theory*. Oxford University Press, New York.

Nau, R. F. 1995. Coherent decision analysis with inseparable probabilities and utilities. *J. Risk Uncertainty* **10** 71–91.

——. 2001. De Finetti was right: Probability does not exist. *Theory Decision* **59** 89–124.

——. 2002. Uncertainty aversion with second-order utilities and probabilities. Working paper, Fuqua School of Business, Duke University, Durham, NC.

——, K. F. McCardle. 1991. Arbitrage, rationality, and equilibrium. *Theory Decision* **31** 199–240.

Pratt, J. W. 1964. Risk aversion in the small and in the large. *Econometrica* **32** 122–136.

——. 1988. Aversion to one risk in the presence of others. *J. Risk Uncertainty* **1** 395–413.

Ross, S. 1976. Risk, return, and arbitrage. I. Friend, J. Bicksler, eds. *Risk and Return in Finance*. Ballinger, Cambridge, MA.

——. 1978. A simple approach to the valuation of risky streams. *J. Bus.* **51**(3) 453–475.

——. 1981. Some stronger measures of risk aversion in the small and the large with applications. *Econometrica* **49** 621–638.

Rothschild, M., J. E. Stiglitz. 1970. Increasing risk I: A definition. *J. Econom. Theory* **2** 225–243.

Rubinstein, M. 1976. The valuation of uncertain income streams and the pricing of options. *Bell J. Econom. Management Sci.* **7** 407–425.

Savage, L. 1954. *The Foundations of Statistics*. John Wiley & Sons, New York.

Schervish, M. J., T. Seidenfeld, J. B. Kadane. 1990. State-dependent utilities. *J. Amer. Statist. Assoc.* **85** 840–847.

Schlee, E., H. Schlesinger. 1993. The valuation of contingent claim markets. *J. Risk Uncertainty* **6** 19–31.

Segal, U., A. Spivak. 1990. First order versus second order risk aversion. *J. Econom. Theory* **51** 111–125.

Shafer, G. 1986. Savage revisited. *Statist. Sci.* **1** 463–501.

Smith, J. E., R. F. Nau. 1995. Valuing risky projects: Option pricing theory and decision analysis. *Management Sci.* **41** 795–816.

Trigeorgis, L., S. P. Mason. 1987. Valuing managerial flexibility. *Midland Corporate Finance J.* **5** 14–21.

Yaari, M. 1969. Some remarks on measures of risk aversion and their uses. *J. Econom. Theory* **1** 315–329.